Analyzing and Visualizing Scalar Fields on Graphs

Yang Zhang
Department of Computer Science and Engineering
The Ohio State University
Columbus, Ohio, USA, 43210
Email: zhang.863@osu.edu

Yusu Wang
Department of Computer Science and Engineering
The Ohio State University
Columbus, Ohio, USA, 43210
Email: yusu@cse.ohio-state.edu

Srinivasan Parthasarathy
Department of Computer Science and Engineering
The Ohio State University
Columbus, Ohio, USA, 43210
Email: srini@cse.ohio-state.edu

arXiv:1702.03825v1 [cs.DB] 10 Feb 2017

***Abstract*—The value proposition of a dataset often resides in the implicit interconnections or explicit relationships (patterns) among individual entities, and is often modeled as a graph. Effective visualization of such graphs can lead to key insights uncovering such value. In this article we propose a visualization method to explore graphs with numerical attributes associated with nodes (or edges) – referred to as scalar graphs. Such numerical attributes can represent raw content information, similarities, or derived information reflecting important network measures such as triangle density and centrality. The proposed visualization strategy seeks to simultaneously uncover the relationship between attribute values and graph topology, and relies on transforming the network to generate a terrain map. A key objective here is to ensure that the terrain map reveals the overall distribution of components-of-interest (e.g. dense subgraphs, k-cores) and the relationships among them while being sensitive to the attribute values over the graph. We also design extensions that can capture the relationship across multiple numerical attributes (scalars). We demonstrate the efficacy of our method on several real-world data science tasks while scaling to large graphs with millions of nodes.**

I. Introduction

Our ability to produce and store data has far outstripped our ability to analyze and utilize this data to derive actionable insight. Many phenomena and problems from all walks of human endeavor can often be represented in graph or network form, where nodes represent entities-of-interest and edges represent interactions or relationships among them. Examples abound across the physical, biological, business, technological, sociological and health domains. A fundamental challenge is the ability to visualize such interconnections at scale while working within the pixel limits of modern displays. This in turn has led to several recent advances in the database systems [1], [2], [3], [4], network science[5], [6], geometry[7], [8], [9], [10], and information visualization [11], [12] communities.

However, the data scientist is often interested in uncovering patterns that go beyond layout and visualization designs – for instance accounting for measures or attributes defined on both nodes and edges of the graph [13], [2]. Such measures often encode information about connectivity – both locally (e.g. triangle density, clustering coefficient, K-Core, K-Truss) as well as globally (e.g. various centrality measures such as closeness, betweenness and harmonic, and also measures of importance such as PageRank and Influence). Beyond just measures from the topology, such measures may also incorporate heterogeneous information related to content (e.g. sequence information of a protein etc.). Visualizing the measure information in such graphs (where each node or edge has one or more scalar measures associated with it) further exacerbates the challenge.

In this article we propose a novel visualization method to explore graphs with numerical measures associated with nodes (or edges) – referred to as *scalar graphs*, and the measure values are referred to as *scalar values*. Each scalar value could either be a natural attribute or a derived attribute summarizing information from multiple natural attributes (e.g. triangle density, centralities, cliquishness) [4], [1]. We model the problem as a continuous function $f : X \to \mathcal{R}$: the domain X is a simplicial complex whose vertex set is the set of input graph nodes, and its topology is determined by the input graph topology. We call such a representation of a graph a *graph-induced scalar field*. We then leverage a powerful "terrain metaphor" idea for visualizing scalar fields [9], [10] and adapt it to our context. Our visualization model *naturally encodes both topological and numerical measure information together*, and is capable of handling large graphs with millions of nodes/edges. Empirically we demonstrate the use of our methodology on a range of data science tasks while providing a comparative assessment against state-of-the-art alternatives from the perspectives of efficacy, efficiency and usability.

Figure 1 previews our methodology for the tasks of: i) visualizing dense subgraphs (K-Cores, K-Trusses, etc); and ii) visualizing communities in social networks. In Figure 1(a), we use K-Core number [5] as a measure on each node and use it to visualize K-Cores in a collaboration network (a fundamental network science task), where the top part of high peaks contains dense K-Cores (Clicking on the circled part of the high peak will show the details of a dense K-Core in the red box). In Figure 1(b), we use community score [14] as a measure to visualize the four communities in a DBLP network, where each major peak is a community, and the sub-peaks in a major peak indicate sub-communities. The top part of a peak contains key members of the community. Our proposed tool can also color the terrain using a second measure, so one can analyze the correlation between two different measures on the graph and furthermore allows for various rotation/transformation operators which allow an end-user to interrogate the data from multiple perspectives. Further details of our evaluation is detailed in Section III. In summary, our contributions are:

- We advocate a novel visualization method to visualize a graph whose vertices/edges are associated with numerical measures. Central to our approach is to embed the graph in *terrain* space. The visualization method enables users to quickly capture how the numerical measure value is distributed over the graph.
- We propose the concept of maximal α-connected component to represent various graph patterns (dense subgraphs, communities, K-Cores, K-Truss etc.), and show that our visualization method not only captures the distribution of graph patterns, but also their global

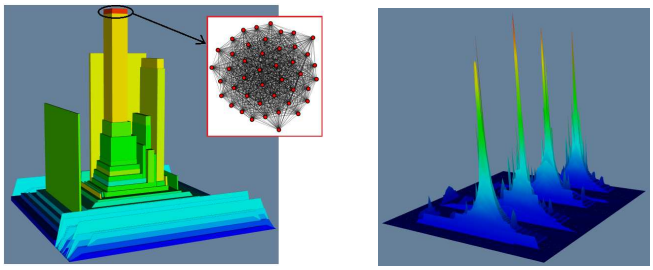

(a) Visualizing K-Cores (b) Visualizing Communities

Fig. 1. Examples of Terrain Visualization

relationships.

- Based on our visualization method, we propose methods to analyze the relation between multiple numerical measures on one graph.
- Finally, we empirically demonstrate that the visualization method is both general-purpose and scalable and can handle graphs with millions of nodes and edges. A key part of our evaluation strategy includes a user-study following best practices (Section IV).

II. Scalar Graph Visualization

Notation: We define a vertex-based scalar graph *G(V, E)* as a graph comprising edges *E* and vertices *V*, where each vertex v has one scalar value associated with it, denoted as *v.scalar*. In the following we refer to vertex-based scalar graph as a *scalar graph* for expository simplicity. We also assume the following notion in the rest of this section. For two subgraphs $G'(V', E')$ and $G''(V'', E'')$ of scalar graph G, G' is the same as G'' (denoted as $G' = G''$) if $V' = V''$ and $E' = E''$, G' is a subgraph of G'' (denoted as $G' \subseteq G''$) if $V' \subseteq V''$ and $E' \subseteq E''$, G' is connected to G'' (denoted as $G' \leftrightarrow G''$) if there is a vertex $v' \in V'$ and a vertex $v'' \in V''$ such that v' is connected to v''.

A. Theoretical Insights

We now describe the key insights underpinning our visualization strategy. We first define **maximal α-connected component** as follows:

Definition 1: A connected component $C(V_C, E_C)$ of scalar graph $G(V, E)$ is a **maximal α-connected component** if it satisfies following conditions:
(1) for every vertex $v \in V_C$, $v.scalar >= \alpha$.
(2) for any vertex $v \in V_C$, if v' is connected to v and $v' \notin V_C$, then $v'.scalar < \alpha$.
(3) for any edge $e(v_1, v_2) \in E$, if $v_1 \in V_C$ and $v_2 \in V_C$, then $e(v_1, v_2) \in E_C$.

Maximal α-connected components are an important aspect of our visualization strategy. We note that the definition accommodates both connectivity (topology) and scalar value information. The scalar values may be inherited from the domain or derived (e.g. local density or k-core values[5])). The distribution of maximal α-connected components in a graph would allow the data scientist to interrogate the distribution and inter-connectivity of such structures within the graph (e.g. distribution and relationships of K-cores).

Definition 2: For vertex v, we define **MCC(v)** as the maximal $v.scalar$-connected component containing v (i.e. $\alpha = v.scalar$).

Theorem 1: For any maximal α-connected component C in G, there is a vertex v in G such that $MCC(v) = C$.

Proof Sketch: It is easy to show that v in Theorem 1 is the vertex with minimum scalar value in C. The proof trivially follows from this observation.

Theorem 2: For two vertices v and v', if $v.scalar = v'.scalar$ and $MCC(v)$ contains v', then $MCC(v) = MCC(v')$.

Proof Sketch: For every vertex v_i in $MCC(v)$, we can show that in $MCC(v)$ there is a path $v_i \rightarrow ... \rightarrow v \rightarrow ... \rightarrow v'$, which starts with v_i, passes through v, and ends at v', and all the vertices on the path have scalar value greater than or equal to $v'.scalar$, so v_i is in $MCC(v')$. Similarly, we can show for every vertex v_j in $MCC(v')$, v_j is in $MCC(v)$. Thus $MCC(v) = MCC(v')$.

Theorem 3: For a maximal α_1-connected component C_1 and another maximal α_2-connected component C_2, if $C_1 \leftrightarrow C_2$ then $C_1 \subseteq C_2$ or $C_2 \subseteq C_1$.

Proof Sketch: Based on Theorem 1, there are two vertices v_1 and v_2 such that $MCC(v_1) = C_1$ and $MCC(v_2) = C_2$. Assume $v_1.scalar \leq v_2.scalar$ (w.l.g). Since $MCC(v_1) \leftrightarrow MCC(v_2)$, for any vertex v' in $MCC(v_2)$, we can find a path $v_1 \rightarrow ... \rightarrow v'$ connecting v_1 and v', and all the vertices on the path have scalar value greater than or equal to $v_1.scalar$, so v' is in $MCC(v_1)$. Thus $MCC(v_2) \subseteq MCC(v_1)$.

Theorems 1 and 2 bound the number of distinct maximal α-connected components that need to be processed by any algorithm (number of vertices in the graph). Theorem 3 highlights an important hierarchical relationship between different maximal α-connected components which will be exploited by our visualization strategy.

B. Vertex Scalar Tree

In this section, we describe the vertex scalar tree (*scalar tree* for short) to analyze a scalar graph. We note that if one views the graph as a 1-dimensional simplicial complex, the vertex scalar values induce a piecewise-linear function on this domain, then the maximal α-connected components we defined previously are akin to level sets or contours for each α [15], [7]. We also note that this perspective also helps us easily extend the notion of scalar tree for edge-based scalar graphs (in Section 2.3).

A scalar tree is a tree in which every node is associated with a scalar value, and the scalar tree T of scalar graph G has the following properties:

1) Every node in T corresponds to a vertex in G with the same scalar value, and vice versa (i.e. one-to-one correspondence).
2) Every maximal α-connected component in G corresponds to a subtree in T, and vice versa (i.e. one-to-one correspondence).
3) Assume a maximal α_1-connected component C_1 corresponds to subtree T_1 in T, and another maximal α_2-connected component C_2 corresponds to subtree

T_2 in T, then C_1 is a subgraph of C_2 if and only if T_1 is subtree of T_2.

4) Assume a maximal α_1-connected component C_1 corresponds to subtree T_1 in T, and another maximal α_2-connected component C_2 corresponds to subtree T_2 in T, then C_1 and C_2 are not connected if and only if T_1 and T_2 are not connected.

Notation: In the following text, the scalar tree node corresponding to vertex *v* is denoted as *n(v)*, and the vertex corresponding to scalar tree node *n* is denoted as *v(n)*. The parent of tree node n is denoted as $parent(n)$. The subtree corresponding to a maximal α-connected component *C* is denoted as *ST(C)*. A subtree *ST* containing nodes $n_1, n_2...n_k$ is denoted as $ST(n_1, n_2, ..., n_k)$, and a connected component C containing vertices $v_1, v_2, ...v_k$ is denoted as $C(v_1, v_2, ..., v_k)$.

Example: In Figure 2 we use an example to illustrate the properties of scalar tree. Figure 2(a) is a scalar graph G, in which every vertex's label is its scalar value. Figure 2(b) is a correspondent scalar tree T of Figure 2(a), rooted at node n_9. Node n_i in Figure 2(b) corresponds to vertex v_i in Figure 2(a) (Property 1). In Figure 2(c), we extract all the maximal 2.5-connected components of G:
$C_1(v_1, v_2, v_3, v_5)$ and $C_2(v_4, v_6)$, and in Figure 2(b), their correspondent subtrees are: $ST(C_1) = ST(n_1, n_2, n_3, n_5)$ and $ST(C_2) = ST(n_4, n_6)$, this satisfies Property 2. We notice that C_1 and C_2 are not connected, and $ST(C_1)$ and $ST(C_2)$ are not connected either, this satisfies Property 4. We also observe that C_1 is a subgraph of a maximal 2-connected component $C_3(v_1, v_2, v_3, v_4, v_5, v_6, v_7)$, and $ST(C_1)$ is a subtree of $ST(C_3) = ST(n_1, n_2, n_3, n_4, n_5, n_6, n_7)$, this satisfies Property 3.

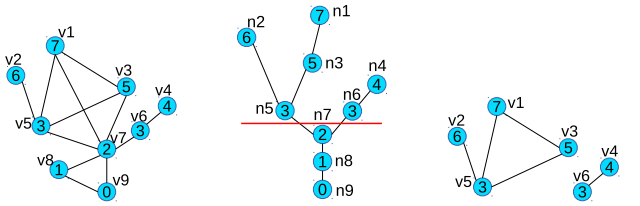

(a) Scalar Graph G (b) Scalar Tree T (c) maximal 2.5-connected components

Fig. 2. Scalar Graph and Scalar Tree

From the properties above, we can see that scalar tree captures the distribution and relationship of all maximal α-connected components within a scalar graph. This allows one to analyze the scalar tree instead of the scalar graph and leverage it for our visualization strategy.

Constructing the Vertex Scalar Tree: Algorithm 1 constructs the scalar tree by leveraging the observation connecting our problem with level sets and contour trees[15], [7]. Algorithm 1 processes all the vertices in decreasing order of scalar values (line 1-3). If the current vertex v_i is connected to a previously processed vertex v_j, but $n(v_i)$ is not in the current subtree of $n(v_j)$ (line 4-5), then we connect $n(v_i)$ with the root of the current subtree of $n(v_j)$ (line 6). Here $root(n(v_j))$ denotes the root node of the current subtree containing $n(v_j)$. Note that now $n(v_i)$ is parent of $root(n(v_j))$, so $n(v_i)$ becomes the new root of the current subtree containing $n(v_j)$.

The running time of line 1 is $O(|V| * log|V|)$. An efficient implementation of line 5 uses the "Union Find" data structure, which compares $root(n(v_i))$ and $root(n(v_j))$. The amortized time for "Union Find" (line 5) is $O(\alpha(n))$ per operation, where

Algorithm 1 Constructing Scalar Tree

Require: A scalar graph G(V, E).
Ensure: The scalar tree T of G.
1: Sort vertices in decreasing order of scalar values, the sorted vertices are $v_1, v_2, ...v_n$;
2: Create a tree node $n(v_i)$ for each vertex v_i;
3: **for** $i = 1$ to n **do**
4: **for** every neighbor v_j of v_i **do**
5: **if** $j < i$ and currently $n(v_i)$ and $n(v_j)$ are not in the same subtree **then**
6: Connect $n(v_i)$ to $root(n(v_j))$; //$n(v_i)$ is parent

$\alpha(n)$ is inverse of Ackermann function, and is usually a small constant. So the total worst-case time cost of Algorithm 1 is $O(|E| * \alpha(n) + |V| * log|V|)$

In the scalar tree T generated by Algorithm 1, every node's scalar value is greater than or equal to its parent's scalar value. If we layout a scalar tree T in such a way that the height of each node is the scalar value of the node, then we could get all the maximal α-connected components for a particular α in a simple way: draw a line with $height = \alpha$ to cross T, and each of the subtrees above the line corresponds to a maximal α-connected component. For example, as Figure 2(b) shows, the two subtrees above red line $height = 2.5$ correspond to all the maximal 2.5-connected components. When every vertex in the input scalar graph G has a distinct scalar value, the scalar tree T generated by Algorithm 1 has the following property:

Proposition 1: For every vertex v in G, assume it corresponds to n(v) in T, then the subtree rooted at n(v) (denoted as ST) in T corresponds to the MCC(v).

Proof: Obviously ST corresponds to a connected component SG. Since $n(v).scalar$ is the minimum scalar value in ST, every vertex in SG has scalar value greater than or equal to $v.scalar$. If there is a vertex v_i that connects to a vertex v_j in SG, and $v_i.scalar > v.scalar$, then due to Algorithm 1, $n(v_i)$ and $n(v_j)$ will be in the same subtree ST, which indicates that v_i is in SG. So SG is a maximal $v.scalar$-connected component, and it is MCC(v). ■

Based on Theorem 1, Theorem 3 and Proposition 1, it follows that when every vertex in the input scalar graph has a *distinct value*, the tree generated by Algorithm 1 has the four properties of the scalar tree. When some vertices in scalar graph have the same scalar value, Algorithm 1 may not guarantee Property 2, and we need to do some additional postprocessing, described next.

Postprocessing the Vertex Scalar Tree: If some vertices in G have the same scalar values, the scalar tree generated by Algorithm 1 may not satisfy Property 2 – a subtree may not correspond to a maximal α-connected component. For example, Figure 3(a) is a scalar graph, and Figure 3(b) is the scalar tree generated by Algorithm 1, in which n_i corresponds to v_i in the scalar graph. The subtree rooted at n_3 is $ST(n_1, n_3)$ corresponds to connected component $C(v_1, v_3)$, however, $C(v_1, v_3)$ is not a maximal α-connected component.

To solve the problem, we use Algorithm 2 to postprocess the scalar tree T generated by Algorithm 1: we merge the ancestor node with all its descendants with the same scalar value into a super node. The correctness of the algorithm is based on the following proposition:

Proposition 2: For any tree node n in T, assume it has an ancestor n_{anc} in T, such that $n_{anc}.scalar = n.scalar$, and $parent(n_{anc})$ is null or $parent(n_{anc}).scalar < n_{anc}.scalar$,

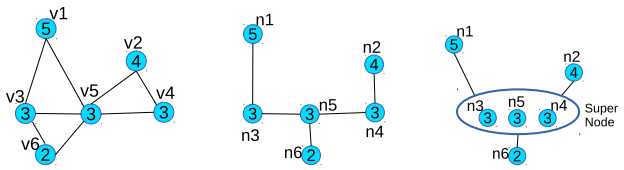

(a) Scalar Graph G (b) Scalar Tree T (c) Super Scalar Tree

Fig. 3. Postprocessing Scalar Tree

then the subtree rooted at n_{anc} (denoted as ST_{anc}) corresponds to the MCC(v(n)).

In the example in Figure 3(b), Algorithm 2 will merge the nodes n_3, n_4, n_5 into a super node, and build a super tree (Figure 3(c)). Every subtree of the super tree will correspond to a maximal α-connected component. After postprocessing, the scalar tree will still satisfy Property 2, 3, 4, but may not satisfy Property 1, because a super node may correspond to multiple vertices, however, this does not affect the further analysis.

Algorithm 2 only needs one pass of the scalar tree T, so the time complexity is $O(|V|)$, where $|V|$ is the number of vertices in the original scalar graph.

Algorithm 2 Postprocess Scalar Tree

Require: The scalar tree T generated by Algorithm 1.
Ensure: The super scalar tree T_{super}.

```
1: Create an array ancestors;
2: Assume the root of T is n_r, n_r.super_node = new Node();
3: ancestors.add(n_r);
4: T_super.add(n_r.super_node);
5: for each node n_anc in ancestors do
6:    Create a queue Q;
7:    Q.push(n_anc);
8:    while !Q.empty() do
9:       n_q = Q.pop();
10:      n_anc.super_node.members.add(n_q);
11:      for every child n_c of n_q do
12:         if n_c.scalar = n_q.scalar then
13:            Q.push(n_c);
14:         else
15:            n_c.super_node = new Node();
16:            ancestors.add(n_c);
17:            T_super.add(n_c.super_node);
18:            Connect n_anc.super_node with n_c.super_node;
```

C. Edge-based Scalar Graph

Here we describe a *novel* approach for modeling scalar values on edges. We define an edge-based scalar graph $G(V, E)$ as a graph comprising edges E and vertices V, where each edge *e* has one scalar value associated with it, denoted as *e.scalar*. Similarly, the **maximal α-edge connected component** is defined below:

Definition 3: A connected component $C(V_C, E_C)$ of scalar graph $G(V, E)$ is a **maximal α-edge connected component** if it satisfies following conditions:
(1) for every edge $e \in E_C$, $e.scalar >= \alpha$.
(2) for any edge $e \in E_C$, if edge e' shares a common vertex with e and $e' \notin E_C$, then $e'.scalar < \alpha$.
(3) for any edge $e(v_1, v_2) \in E_C$, we have $v_1 \in V_C$ and $v_2 \in V_C$.

The edge scalar tree T of edge scalar graph G has the following properties:

1) Every node in T corresponds to an edge in G with the same scalar value, and vice versa (i.e. one-to-one correspondence).
2) Every maximal α-edge connected component in G corresponds to a subtree in T, and vice versa (i.e. one-to-one correspondence).
3) Assume a maximal α_1-edge connected component C_1 corresponds to subtree T_1 in T, and another maximal α_2-edge connected component C_2 corresponds to subtree T_2 in T, then C_1 is a subgraph of C_2 if and only if T_1 is subtree of T_2.
4) Assume a maximal α_1-edge connected component C_1 corresponds to subtree T_1 in T, and another maximal α_2-edge connected component C_2 corresponds to subtree T_2 in T, then C_1 and C_2 are not connected if and only if T_1 and T_2 are not connected.

Naive Method: a naive way to build edge scalar tree is to first convert the edge-based scalar graph $G(V, E)$ to a dual graph $G_d(V_d, E_d)$– where every edge in G is converted to be a vertex in G_d, and if two edges in G share a common vertex, their correspondent vertices in G_d are connected. We then apply Algorithm 1 to G_d – the generated tree is an edge scalar tree of G. The time complexity of the method is $O(|E_d| * log|V_d| + |V_d| * log|V_d|)$. In the dual graph G_d, we have $|V_d| = |E|$ and $|E_d| = O(\sum_{v \in V} degree(v)^2)$, so the time complexity is actually $O(\sum_{v \in V} degree(v)^2 * log|E| + |E| * log|E|)$. The time cost is high because the bottleneck $\sum_{v \in V} degree(v)^2$ could be $|V|^3$ in the worst case.

Optimized Method: We propose a novel, more efficient method (Algorithm 3) to construct edge scalar tree from the edge scalar graph, and the time complexity is reduced to be $O(|E| * log|E|)$. In line 1, we sorted all the edges in descending order of scalar value. In line 2-3, we select the min_id_edge on vertex v that has the minimum index. In line 6-9, we process edge e_i. Instead of checking all e_i's neighbor edges (edges which share a common vertex with e_i) , we just need to check the min_id_edges of e_i's two vertices (line 6-7). This is based on Proposition 3:

Proposition 3: If edge e_j is a neighbor edge of e_i ($i > j$), and they share the same vertex v, when processing e_i in line 6-9 of Algorithm 3, root($n(e_j)$) is the same as root($n(v.min_id_edge)$).

Proof: Since $i > j$, e_j will be processed before e_i. When processing e_j in line 6-9 of Algorithm 3, $n(e_j)$ will be connected to root($n(v.min_id_edge)$), which means $n(e_j)$ and $n(v.min_id_edge)$ will be in the same subtree thereafter. So when processing e_i, root($n(e_j)$) is the same as root($n(v.min_id_edge)$). ■

The time complexity of line 1 in Algorithm 3 is $O(|E| * log|E|)$. For each edge e, line 8 is executed $O(1)$ times, so line 8 is executed a total of $O(|E|)$ times, and the total running time of line 5-9 is $O(|E| * log|E|$. The worst case running time of Algorithm 3 is $O(|E| * log|E|)$. Similar to the case described in the previous section, if some edges have the same scalar value, we can use Algorithm 2 to postprocess the edge scalar tree.

D. Relationship between maximal α-(edge) connected component and Dense Subgraph

A dense subgraph is a connected subgraph in which every vertex is heavily connected to other vertices in the subgraph.

Algorithm 3 Constructing Edge Scalar Tree

Require: An edge scalar graph G(V, E).
Ensure: The edge scalar tree T of G.
1: Sort edges in decreasing order of scalar values, the sorted edges are $e_1, e_2, ...e_n$.
2: **for** each vertex v in G **do**
3: set $v.min_id_edge$ to be the edge on v that has the minimum index.
4: Create a tree node $n(e)$ for each edge e.
5: **for** $i = 1$ to n **do**
6: Assume e_i has two vertices v_1 and v_2, create an array: $min_neighbors[2] = \{v_1.min_id_edge, v_2.min_id_edge\}$;
7: **for** each edge e_m in $min_neighbors$ **do**
8: **if** $m < i$ and currently $n(e_i)$ and $n(e_m)$ are not in the same subtree **then**
9: connect $n(e_i)$ to root($n(e_m)$) //$n(e_i)$ is parent

K-Core [5] and K-Truss [16], [17](also called Triangle K-Core in [4], DN-graph in [3]) are two common dense subgraph patterns that draw much attention in recent works. The definitions of K-Core and K-Truss are as follows:

Definition 4: A K-Core is a subgraph in which each vertex participates in at least K edges within the subgraph. The maximal K-Core of a vertex v is the K-Core containing v that has the maximum K value, and the K value of maximal K-Core of v is denoted as KC(v).

Definition 5: A K-Truss is a subgraph in which each edge participates in at least K triangles within the subgraph. The maximal K-Truss of an edge e is the K-Truss containing e that has the maximum K value, and the K value of maximal K-Truss of e is denoted as KT(e).

Now we prove the relationship between maximal α-(edge) connected component and K-Core, K-Truss below.

Proposition 4: In a scalar graph G where for any vertex *v*, *v.scalar = KC(v)*, a maximal α-connected component in G is a K-Core where $K = \alpha$.

Proof: Assume in a maximal α-connected component C, vertex v has the minimum scalar value. Based on the definition of K-Core, for every vertex *v'* in the maximal K-Core of *v*, $KC(v') >= KC(v)$, so the maximal K-Core of *v* is a subgraph of *C*. Since *v* is connected to at least $KC(v)$ vertices in its maximal K-Core, *v* is connected to at least $KC(v)$ vertices in *C*. For every other vertex *v'* in *C*, similarly we can get that *v'* is also connected to at least $KC(v)$ vertices in *C*. So *C* is a K-Core where $K = KC(v)$, since $KC(v) = v.scalar >= \alpha$, C is also a K-Core where $K = \alpha$. ■

Proposition 5: In an edge scalar graph G where for any edge *e*, $e.scalar = KT(e)$, a maximal α-edge connected component in G is a K-Truss where $K = \alpha$.

The proof is similar to the proof of Proposition 4 and is omitted in the interests of space. Note that when we define the scalar value of each vertex/edge to be KC(v)/KT(e), the (edge) scalar tree will capture the distribution and relationships among K-Cores or K-Trusses in the graph.

E. Visualization via Terrain Metaphor

Scalar trees are usually not easy to visually interpret, especially when the size of the tree is too large. We adapt the terrain metaphor – topological landscape visualization technique defined on scalar-valued functions [10] – to visualize scalar graphs.

In Figure 4 we use an example to illustrate how to convert the scalar tree in Figure 4(a) to terrain visualization in Figure 4(c). In Figure 4(b) we first layout all the tree nodes of Figure 4(a) in a 2D plane, every node n_i is represented by a boundary b_i in the 2D plane, and the area enclosed by the boundary b_i is proportional to the number of nodes in subtree (not including n_i) rooted at n_i. To generate the 2D layout, we start traversing the tree from the root(bottom) node n_9, draw the outermost boundary b_9 to represent it. Then we move to n_8, and draw a boundary b_8 inside b_9. When we reach n_7, and draw boundary b_7, we find there are two subtrees rooted at n_7, so we split the area inside b_7 into 2 areas, and recursively layout each subtree in each area. When we reach leaf nodes n_1, n_2, n_4, since the size of their subtrees is 0, their correspondent boundaries degenerate to be points.

To convert the 2D layout (Figure 4(b)) into a terrain visualization in 3D space (Figure 4(c)), we first escalate each boundary b_i in Figure 4(b) to the height of $n_i.scalar$, and then draw a "wall" between neighboring boundaries. Finally we generate a terrain in Figure 4(c). We can color the terrain by assigning each vertex a color value, and since each "wall" is confined by two boundaries b_i and b_j, we color the wall based on the color value of the vertex corresponding to b_i or b_j. [1]

To identify a subtree of the scalar tree in the terrain visualization, we locate the correspondent boundary b_r of the subtree root n_r, and the terrain area within the boundary b_r corresponds to the subtree. In our paper, we define a $peak_\alpha$ in terrain as below:

Definition 6: A $peak_\alpha$ is the terrain area within a boundary whose height is α.

Since each $peak_\alpha$ corresponds to a subtree in scalar tree, we can easily get that every $peak_\alpha$ corresponds to a maximal α-connected component. Also, $peak_\alpha$s preserve the containment/connection relationship of maximal α-connected components. For example, the red peak in Figure 4(f) is a $peak_5$, which corresponds to the maximal 5-connected component (red nodes) in Figure 4(d), and the red peak in Figure 4(i) is a $peak_3$, which corresponds to the maximal 3-connected component (red nodes) in Figure 4(g). One $peak_\alpha$ may contain some sub-peaks, which indicates its maximal α-connected component contains other maximal α'-connected components. For example, $peak_5$ in Figure 4(f) is contained in $peak_3$ in Figure 4(i), this indicates that the correspondent maximal 5-connected component in Figure 4(d) is a subgraph of the maximal 3-connected component in Figure 4(g). In a $peak_\alpha$, the area of its bottom boundary indicates the number of vertices in its correspondent maximal α-connected component.

To get all the maximal α-connected components for a particular α value, we can use a 2D plane with height = α to cross the terrain in 3D space, and all the $peak_\alpha$s above the plane correspond to all maximal α-connected components. *The benefit of using terrain visualization is, it captures the overall information of all maximal α-connected components in one picture. Also, we could encode more information in the terrain by using colors to the terrain.*

[1] The use of 3D rather than 2D was a conscious one. First we found that the 3D abstraction better matched users' mental-map of the "terrain" concept as well as the hierarchical relationships amongst components-of-interest. While 3D-layouts potentially pose a problem with respect to occlusion, the ability to interactively rotate the point-of-view along with the ability to link 2D-layouts of regions-of-interest as discussed below, alleviates this issue.

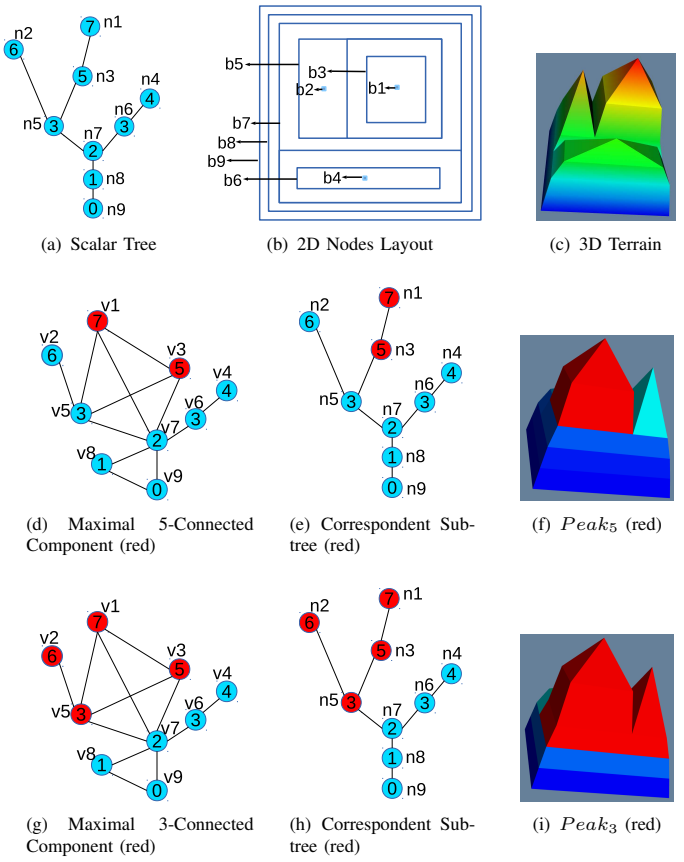

(a) Scalar Tree (b) 2D Nodes Layout (c) 3D Terrain

(d) Maximal 5-Connected Component (red) (e) Correspondent Subtree (red) (f) $Peak_5$ (red)

(g) Maximal 3-Connected Component (red) (h) Correspondent Subtree (red) (i) $Peak_3$ (red)

Fig. 4. Terrain Visualization of a Simple Scalar Tree

User Interaction: Our terrain visualization tool provides following features to help users interact with the terrain.

Rotate: the user could rotate the terrain to look at it from different angles. For example, Figure 4(c) and Figure 4(f) show the same terrain from two different viewpoints.

Zoom in/out: the user can zoom in/out to see the details/overview of the terrain. For example, in Figure 8(a), we zoom into the terrain in the left picture, and get a clear picture of the two peaks in the right picture.

Simplification: When visualizing a scalar tree with too many nodes, the rendering and interaction speed might be slow, we simplify the tree to make the visualization faster as follows. We discretize the scalar values, so similar scalar values will be approximately represented by the same value, and then we can use Algorithm 2 to build an approximate super tree with far fewer tree nodes.

Linked-2D-Displays: Our tool allows the user to select any region of the terrain, and invoke a "callback" function to visualize the selected region using other visualization method. For example, in Figure 6(c), we select the region in the white dashed line box, and 2D-linked spring layout visualization method to draw the selected region in the red box beside it. Expert users optionally can use the "callback" function to integrate our terrain visualization with other customized visualization methods. We can also link a 2D treemap of the scalar graph by setting the height of all boundaries to 0 and (optionally) using colors – red/yellow/green/blue – to indicate highest/high/low/lowest value – so the red/yellow blocks in the 2D treemap indicate the subgraph areas with high scalar values (see Figure 5(a) and Figure 5(b)). The 2D visualization clearly shows how the red/yellow/green blocks are distributed over the graph while in the 3D visualization the user will need to rotate her point-of-view to get a similar map. That said, using color to encode scalar value has a drawback. For example, peak 1 and peak 2 in Figure 5(b) correspond to block 1 and block 2 in Figure 5(a), from the height we can see that the peak 2 is a little higher than the peak 1, but from the color we cannot tell the difference between block 1 and block 2. Finally, note that using height allows us to color the terrain based on a different attribute (from the one that generates the terrain), which helps to get a picture of the correlation across multiple scalars, as we shall discuss next.

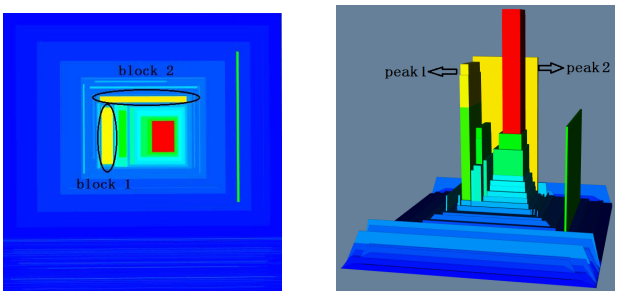

(a) GrQc 2D visualization (b) GrQc 3D visualization

Fig. 5. 2D Treemap vs. 3D Terrain

F. Handling Multiple Scalar Fields

On some graphs, multiple scalar fields can be defined, which means there are multiple scalar values defined on each vertex. For example, a vertex v has degree value and KC(v) value. Users might be interested in how the multiple scalar fields correlate with each other – are their changing trends the same or the opposite on the graph? In this section, we propose two indexes to measure the correlation of two scalar fields on a graph. Then we can use terrain visualization to analyze the relationship between two scalar fields.

Local Correlation Index: Assume we have two scalar fields, S_i and S_j, defined on a graph. Each vertex v has scalar values $v.scalar_i$ and $v.scalar_j$ in the two scalar fields. Some previous work [8] proposed a measure to compute the correlation of multiple scalar fields in continuous domain, we adapt their method and propose Local Correlation Index to measure the correlation of two scalar fields on local areas of a graph. Here local area is defined as k-hop neighborhood of each vertex v(denoted as N(v)), for all experiments we limit this to be 1-hop. The Local Correlation Index of S_i and S_j on N(v) is denoted as $LCI_{Si,Sj}(v)$. For each vertex v, we compute $LCI_{Si,Sj}(v)$ as follows.

$$\overline{v.scalar_i} = \frac{\sum_{u \in N(v)} u.scalar_i}{|N(v)|}$$

$$Cov_{ij}(v) = \frac{\sum_{u \in N(v)} (u.scalar_i - \overline{v.scalar_i}) * (u.scalar_j - \overline{v.scalar_j})}{|N(v)|}$$

$$LCI_{Si,Sj}(v) = \frac{Cov_{ij}(v)}{\sqrt{Cov_{ii}(v)} * \sqrt{Cov_{jj}(v)}}$$

$LCI_{Si,Sj}(v)$ is actually the correlation of the scalar values of S_i and S_j on v's k-hop neighborhood. A positive/negative

$LCI_{Si,Sj}(v)$ indicates the changing trends of S_i and S_j are consistent/inconsistent on v's k-hop neighborhood. This method can easily be adapted to analyze edge-based scalar graphs.

Global Correlation Index: We can compute the Global Correlation Index (GCI) of scalar fields S_i and S_j on a graph by averaging the Local Correlation Indexes of all neighborhoods.

$$GCI_{S_i,S_j}(G) = \sum_{v \in V} LCI_{S_i,S_j}(v)/|V|$$

By comparing the Global Correlation Index and Local Correlation Index, we may identify some outlier neighborhood on which the correlation of scalar fields S_i and S_j is different from the overall correlation.

Terrain Visualization: To visualize the local correlation between two scalar fields, we can use $LCI_{Si,Sj}(v)$ as a scalar field to draw the terrain. This will show us the overall distribution of $LCI_{Si,Sj}(v)$ over the graph, and help us identify the area of the graph where the two scalar fields are positively/negatively correlated.

We can also visually capture the global correlation of scalar fields S_i and S_j through coloring terrain visualization. We use one scalar field S_i to draw the terrain visualization, and use the other scalar field S_j to color the terrain (see Figure 1(a)). Please note that this method can also visualize the relationship between a numerical attribute and a nominal attribute of vertices, by coloring the terrain based on the value of nominal attributes.

G. Related Context

We are now in a position to briefly place the proposed visualization strategy in the context of related work. Visualizing graph data is an important problem (See [18] for a recent survey – due to interests of space here we focus only on the most relevant). Gronemann et al. [19] (similarly Athenstadt et al. [20]) use topographic maps to visualize clustering structure within a graph – each mountain corresponds to a cluster. van Liere et al.[21] propose the GraphSplatting method to visualize a graph as a 2D splat field. Telea et al.[22] generate a concise representation of graph by clustering edges and bundling similar edges together and subsequently visualize the graph. While effective on small scale datasets for displaying overall cluster structure, these methods simply do not scale to large data with millions of edges nor do they account for attributed graphs (scalar values). Bezerianos et al. [23] and Wattenberg [24] propose interactive visual system to let users explore networks with multiple node or edge attributes. However, these methods do not consider hierarchical relationships among components-of-interest and graph attributes simultaneously.

Sariyuce et al. [25] proposed (r, s)-nucleus to denote a dense subgraph comprised by cliques, and use forest of nucleus to represent hierarchical structure of a graph. The difference is, their definition of (r, s)-nucleus focuses on density of subgraph, while our definition of ***maximal α-connected components*** focuses on relation between scalar values and graph topology, and not just limited to nucleus motifs. Moreover, their effort does not consider visualization as an objective – a primary focus of our effort. Some visualization methods (such as LaNet-vi [6]) are proposed to visualize K-Cores, but they are not general enough to handle other vertex/edge attributes. Martin et al. [26] proposed OpenOrd to visualize large scale graphs in multi-level way, but do not effectively highlight components-of-interest. Both LaNet-vi and OpenOrd share some of our objectives w.r.t network visualization, and we compare and contrast with these efforts in Section IV.

In summary, a major difference between our work and prior art is that we propose to analyze the graph through the hierarchical structure (scalar tree) induced by the ***maximal α-connected components***. The benefit is that it *naturally* encodes the relationship between clustering structure and scalar values – it highlights how scalar values evolve from high values to low values over the graph. This is particularly useful for a data scientist who wishes to understand how a community is expanded from its core members to peripheral members(see Figure 8)[27], [28], [29]. Furthermore, based on different attributes, the ***maximal α-connected component*** can represent different subgraph patterns, such as K-core, K-truss, subcommunity, which has attracted much interest within the database community [1], [5], [16], to reveal the topological relationship (containment, connection) among components of interest (e.g. K-Cores, K-Trusses, communities). We examine these issues next.

III. Experimental Evaluation

We seek to evaluate the effectiveness (qualitative) and efficiency of the our interactive network visualization method in this section. We leverage a wide range of datasets from the network science community (some are from [30]) as noted in Table I. Heights in our terrain visualization represent scalar measures of input scalar graph while color represents intensity of the same measure (unless otherwise noted). The color ranges from red (most intense); yellow (intense); green (less intense); blue (least intense). All experiments are evaluated on a 3.4GHz CPU, 16G RAM Linux-based desktop.

TABLE I. Dataset Properties

Dataset	# Nodes	# Edges	Context
GrQc	5242	14496	Coauthorship in General Relativity and Quantum Cosmology
Wikivote	7115	103689	Who-votes-on-whom relationship between Wikipedia users
Wikipedia	1,815,914	34,022,831	Links between Wikipedia pages
PPI	4741	15147	Protein Protein Interaction network
Cit-Patent	3,774,768	16,518,947	Citations made by patents granted between 1975 and 1999
Amazon	334863	925872	Co-Purchase relationship between products in Amazon
Astro	17903	196972	Coauthorship between authors in Astro Physics
DBLP	27199	66832	Coauthorship between authors in (Database, Data Mining, Machine Learning, Information Retrieval)

A. Visualizing Dense Subgraphs

Effectiveness: The visualization of dense subgraphs within graphs has been of much interest within the database and information visualization. Examples abound and include CSV plots [1], K-Core [6] and Triangle K-Core (K-Truss) [4] plots. Here we use our terrain visualization to visualize K-Cores and K-Trusses and compare with the previous methods.

We consider two datasets (GrQc, Wikivote) for this illustration. In Figure 6(a) and Figure 6(b), we use the traditional spring layout Algorithm [31] to draw both networks – it is hard to say anything about the distribution of dense subgraphs using such a plot. Following the discussion in Section II-D, we use KC(v) as scalar value, and generate the terrain visualization of both networks in Figure 6(c) and Figure 6(d). Recall that if the scalar value of every vertex v is defined to be KC(v), each maximal α-connected component is a K-Core where $K = \alpha$.

Thus in the terrain, each $peak_\alpha$ is a K-Core where $K = \alpha$. The distribution of K-Cores in the two datasets is obviously different. Figure 6(d) shows that there is one single high peak, which means the network has one densest K-Core, and K-Core density gradually decreases to the neighboring vertices. Figure 6(c) shows that there are several high peaks within the GrQc network, which means there are several disconnected K-Cores with high K values (dense K-Cores).

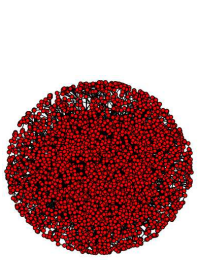

(a) GrQc (spring layout)

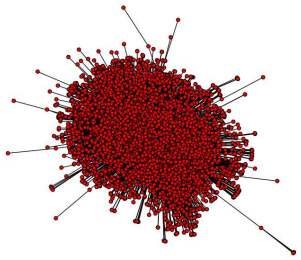

(b) wikiVote (spring layout)

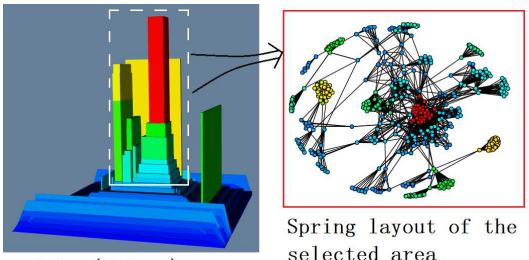

(c) GrQc (K-Core)

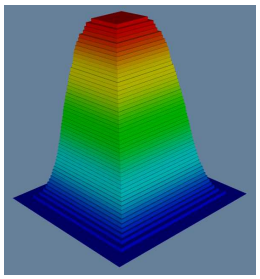

(d) wikiVote (K-Core)

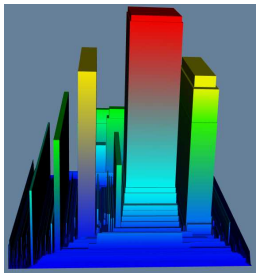

(e) GrQc (K-Truss)

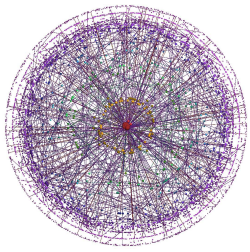

(f) 2D Visualization of K-Cores in GrQc

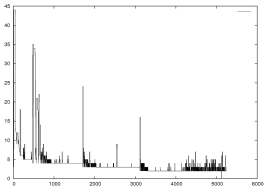

(g) 2D plot of K-Trusses in GrQc

Fig. 6. Visualizing Dense Subgraphs the Network

Moreover, Figure 6(c) clearly illustrates the hierarchical relationship among K-Cores. In the selected terrain area (the terrain area in dashed line), the red peak is placed on green and blue foundation, which means the dense K-Core is contained in some less dense K-Cores. This can be verified by drawing spring layout of the selected region in the red box (with our tool, a user can select a region, and call other visualization methods to draw the selected region), the red dense K-Core is surrounded by some green and blue vertices. The visualization of hierarchy is important, as it allows an analyst to derive high level insights on the connectivity that is not immediately obvious even in state of the art K-Core plots as shown in Figure 6(f) [6] for the GrQc network. We give more detailed comparison between terrain visualization and other visualization methods in User Study section.

Also we can color the terrain using a second measure. In Figure 1(a), we color the terrain based on vertex degree (red/yellow/green/blue area indicates vertices with highest/high/low/lowest degrees), we can see that generally KC(v) is positively correlated with degree – vertices in dense K-Cores have high degrees, except a few outlier vertices that have relatively high degree but low KC(v) values (the yellow area at the bottom of the terrain). They are usually local hub nodes with sparse neighborhood.

We can illustrate the same principle when visualizing K-Trusses (used to understand triangle density) instead of K-Cores. Here each edge uses KT(e) as the scalar measure, and we use the edge-based scalar graph for visualizing the K-Trusses in GrQc dataset. The terrain visualization is in Figure 6(e) where high peaks indicate dense K-Trusses. To contrast, Figure 6(g) depicts a CSV plot, a state-of-the-art density plot leveraged within the database community [1], [4]. Again such visualization strategies do not reveal important hierarchical relationships (e.g. contains) among different K-Trusses. Also we note that our visualization method is a common and flexible framework which can render plots based on different scalar measures, and the ability to rotate, filter and extract details on demand (allowing the analyst to quickly identify regions of interest) will help users understand the graph data better.

Scalability: We next examine the efficiency of Algorithm 1, 2 and 3. Every dataset has duplicate scalar values, so the generated trees are all super trees. We test our methods on datasets of various sizes, and list the number of nodes in the final super (edge) scalar tree(N_t), time cost to construct the tree (tc) and visualize the tree (tv) in Table II. The time cost to construct the tree (tc) includes the time cost to construct the tree (Algorithm 1 or 3) and postprocess the tree (Algorithm 2). The time cost to visualize the tree (tv) is the time cost for the visualization software to read the scalar tree and render the terrain visualization.

TABLE II. TERRAIN VISUALIZATION TIME COST(SEC)

Dataset	Scalar	N_t	tc	te	tv
GrQc	KC(v)	869	0.0018		<1
GrQc	KT(e)	728	0.0039	0.0072	<1
WikiVote	KC(v)	106	0.0037		<1
WikiVote	KT(e)	44	0.053	0.69	<1
Wikipedia	KC(v)	230	6.9		2
Wikipedia	KT(e)	1,903	49.3	16334	22
Cit-Patent	KC(v)	1,059	7.1		2
Cit-Patent	KT(e)	110,412	27.7	65.3	13

Also we list the time cost of the naive method (using dual graph) to build edge-scalar tree (te) in Table II. We can see the improved method (tc) is much faster than the naive method (te), especially on the Wikipedia dataset, the improved method is more than 300 times faster than the naive method.

Additionally, in Figure 7, we show the terrain visualizations of (edge) scalar tree of wikipedia and cit-Patent datasets. The peaks in Figure 7(a) 7(b) 7(c) 7(d) indicate dense K-Cores and K-Trusses in the network, we highlight the highest peaks in

Figure 7(b) and Figure 7(c), and draw the details in Figure 7(e) and Figure 7(f). They are a K-Truss with K = 86 and a K-Core with K = 64.

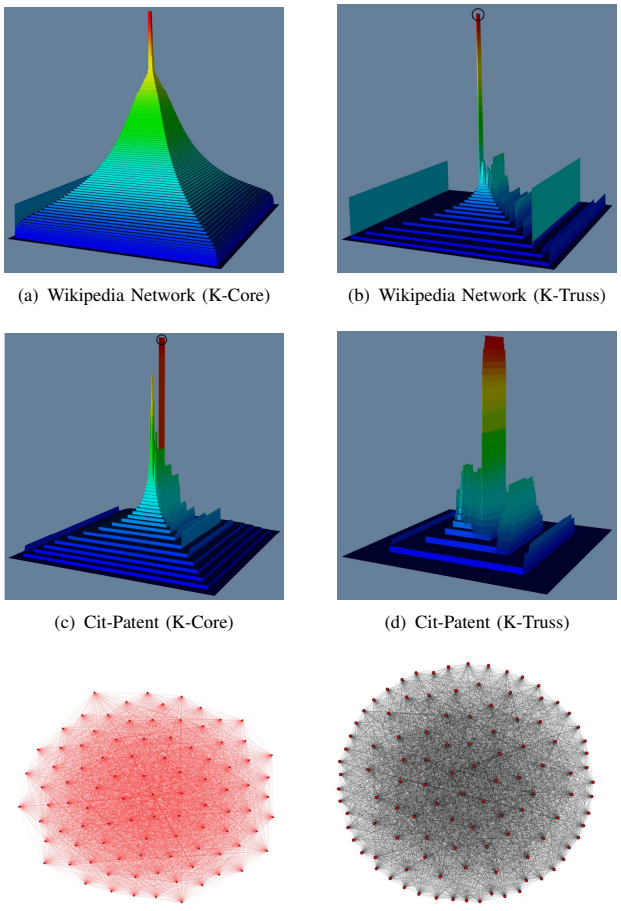

(a) Wikipedia Network (K-Core) (b) Wikipedia Network (K-Truss)

(c) Cit-Patent (K-Core) (d) Cit-Patent (K-Truss)

(e) Densest K-Truss of Wiki (f) Densest K-Core of Cit

Fig. 7. Visualizing K-Cores and K-Trusses

B. Visualizing Communities and Roles

Visualizing Community Affiliation: We next illustrate the flexibility of the terrain visualization scheme on an important network science task – understanding community structures (see Figure 8). We use a subset of the DBLP network (DBLP(sub)) for this purpose, comprising authors who publish in the areas of Machine Learning, Data mining, Databases and Information Retrieval. We apply a state-of-the-art overlapping (soft) community detection algorithm [14] on this dataset to detect four communities. Each author in the dataset is affiliated with a community score vector (c_0, c_1, c_2, c_3) indicating how much it belongs to each community. To visualize the affiliation of a particular community i, we use c_i as the corresponding scalar measure, and draw the terrain of the network. $peak_\alpha$ in the terrain indicates a connected component in which every vertex has $c_i \geq \alpha$.

In Figure 8(a), we visualize community 1, in which most authors are database researchers. We highlight two peaks in the circle of Figure 8(a), and zoom in to get a clear picture of the two peaks on the right. Our tool allows us to easily select authors (vertices) in each peak. We find that authors in the left peak include researchers *Donald Kossmann, Divyakant Agrawal, Amr El Abbadi, Michael Stonebraker, Samuel Madden, and Joseph M. Hellerstein* while authors in the right peak include *Zheng Chen, Hongjun Lu, Jeffrey Xu Yu, Beng Chin Ooi, Kian-Lee Tan, Qiang Yang and Aoying Zhou.* Since authors in both peaks have high community scores (c_1), they can be seen as core members of the community although from different geographic areas. The fact that they are in two separate peaks indicates that authors in one peak do not work with authors in the other peak in the dataset. Similarly, we also observe subcommunities in another community (Figure 8(b)) largely comprising Machine Learning researchers. We also find two peaks in the terrain, and authors in the left peak are *Philip S. Yu, Christos Faloutsos, Michael I. Jordan, Stuart J. Russell, Daphne Koller, Sebastian Thrun, Wei Fan and Andrew Y. Ng,* who all work in United States, while authors in the right peak are *Hang Li, Ji-Rong Wen, Tie-Yan Liu, Lei Zhang, Wei-Ying Ma, Qiang Yang and Yong Yu,* who are researchers in China.

Figure 1(b) visualizes the four communities together to give an overview of them.

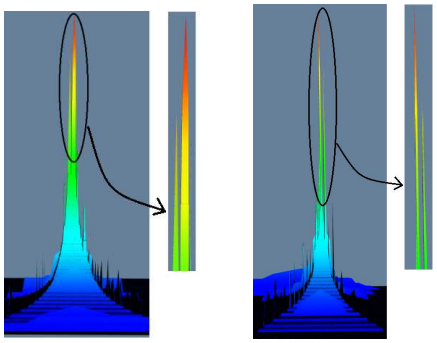

(a) Community 1 (b) Community 2

Fig. 8. Visualizing two communities in DBLP network

Visualizing Role and Community Affiliation: Moving beyond community affiliation, the ability to uncover the roles of individual nodes (e.g. bridge, hub, periphery and whisker) within a network or community has received recent interest [32], [33]. Here we examine how one may use the terrain to visualize the distribution of roles over a community. We leverage a recent idea to simultaneously detect communities and roles on large scale networks [33]. For each vertex in the network the algorithm outputs a community affinity vector $(c_1, \ldots, c_m)$ and a role affinity vector $(r_1, \ldots, r_n)$.

As before we focus on a particular community (community i) and use the community score (c_i) of each vertex to create terrain visualization. The peak in Figure 9(a) contains the vertices affiliated with one major community in Amazon co-purchase network. Instead of re-using the intensity of community score (c_i) to color vertices we actually use the dominant role for each vertex (four roles is typical [32]) to color vertices. We assign each role a color, the "hub vertex" is green, the "dense community vertex" is blue, the "periphery vertex" is red. Then we assign the color of roles to the terrain in Figure 9(a). From the terrain visualization, we can see that the vertices in the community have 3 roles, the hub vertex has the highest community score (green top), and below it is the blue portion, which means the "hub vertex" is surrounded by some "dense community vertices" in the network. The red part of the peak indicates that there are some "peripheral vertices" attached to the community. Since the community contains a small number of vertices, we can draw the details of the community using node-link visualization in Figure 9(b).

All nodes in Figure 9(b) are books on Amazon, and we list a few of them in Table III. The green node is the book which has the highest salesrank and is focused on creativity (hub). Most of the blue nodes are books about creativity (densely

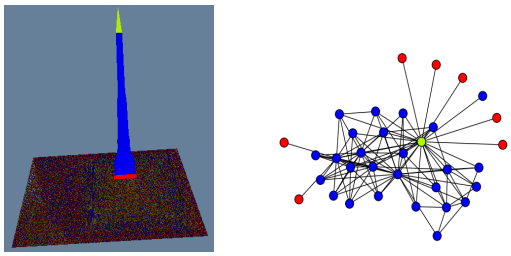

(a) Roles on a community (b) Detail of the community

Fig. 9. one community of Amazon co-purchase network

connected), while the red nodes are books loosely relevant to creativity (periphery).

TABLE III. BOOK NAMES OF SOME NODES IN FIGURE 9(B)

Role	Book Name
green	Artist's Way, The PA
blue	Heart Steps: Prayers and Declarations for a Creative Life
blue	Inspirations: Meditations from the Artist's Way
blue	Reflections on the Artist's Way
blue	The Artist's Way Creativity Kit
red	Writing From the Inner Self
red	Codes of Love : How to Rethink Your Family and Remake Your Life

C. Comparing Different Centralities

In this section we examine the use of our approach for understanding the relationship of two different measures of centrality across various nodes within a network. We will compare two centralities, degree centrality and betweenness centrality, as two scalar fields, S_d and S_b.

We use the Astro Physics collaboration network, in which each author is a vertex, and each edge indicates a coauthorship between two authors. We first compute the Local Correlation Index of each vertex (as described earlier), and then compute the Global Correlation Index of the network, GCI_{S_d,S_b} = 0.89. This indicates that the overall correlation between degree centrality and betweenness centrality is highly positive.

In this case, we are interested in those vertices with negative LCI values, as they could be seen as outliers. We define an outlier score for each vertex v as follows:

$$outlier_score(v) = -LCI(v)$$

the vertex with more negative LCI(v) will have higher outlier score. We use $outlier_score(v)$ as scalar field to draw the terrain in Figure 10(a), and color the terrain using S_d (degree centrality), where red/yellow/blue indicates high/moderate/low degree. We notice that most high peaks are blue, which indicates that the outlier vertices usually have low degree.

We drill down into the two peaks in the terrain within black and red circles, and select the vertex at the top of each peak. Our software allows us to integrate a different visualization method to inspect the selected two vertices. In this case, we use spring layout to draw the two vertices' 2-hop neighborhoods in Figure 10(b) and Figure 10(c). We pick these two vertices specifically because one is in high peak while the other appears to be in a broader but smaller peak. In both cases the vertices picked have high outlier score, which indicates the correlation between degree centrality and betweenness centrality is negative in their neighborhood. Actually the two vertices have relatively higher betweenness and lower degree centrality when compared to many of their

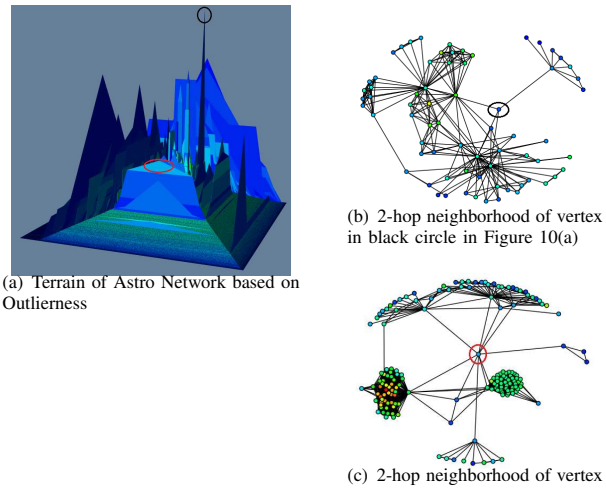

(a) Terrain of Astro Network based on Outlierness

(b) 2-hop neighborhood of vertex in black circle in Figure 10(a)

(c) 2-hop neighborhood of vertex in red circle in Figure 10(a)

Fig. 10. Compare Degree and Betweenness Centralities

neighbors. From Figure 10(b) and Figure 10(c) we can see the two vertices (in the circles) are like bridge nodes connecting multiple communities.

D. Query Result Understanding

Our terrain visualization can also be extended to visualize the results of SQL queries. Here we consider a database of plant genus curated by OSU's Horticulture department. In the interests of space, we focus on a common query posed to this dataset, specified by a domain expert, and model the output (a 5 dimensional materialized table) as a nearest-neighbor (NN) graph (distance measure and threshold again specified by domain expert) and then visualize the graph using the terrain visualization (Figure 11). Color represents different plant genus (3 types in query output), height is a scalar value representing the values of two of the selected attributes from the query result (attributes 1 and 2). While details of the genus are omitted for expository simplicity, the query result visualization clearly conveys the following: i) the result set from the SQL query contains three plant genus (red, green and blue) of which the blue genus is well separated from the other two; ii) It is also clear that the (red) genus is closer to the (green) genus and is in some senses contained within it, i.e. more central, from a connectivity standpoint (within the NN graph); iii) finally attribute 1 demonstrates greater genus separability (variance in terrain heights across genus) on the subset of data produced by this particular query. While preliminary in nature, such visualizations can potentially allow the domain expert to better understand the coherency of the output w.r.t the selection predicates (attributes) of the query.

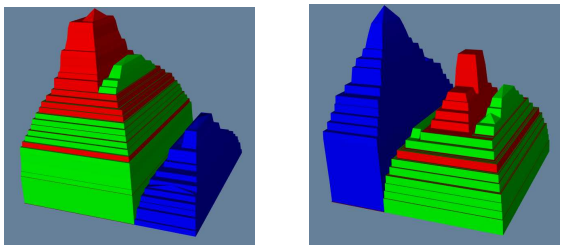

(a) Attribute 1 as Scalar value (b) Attribute 2 as Scalar value

Fig. 11. Terrain Visualization of Plant Dataset Query Result

E. Related Context

We are now in a position to briefly place the proposed visualization strategy in the context of related work. Visualizing graph data is an important problem (See [18] for a recent survey – due to interests of space here we focus only on the most relevant). Gronemann et al. [19] (similarly Athenstadt et al. [20]) use topographic maps to visualize clustering structure within a graph – each mountain corresponds to a cluster. van Liere et al.[21] propose the GraphSplatting method to visualize a graph as a 2D splat field. Telea et al.[22] generate a concise representation of graph by clustering edges and bundling similar edges together and subsequently visualize the graph. While effective on small scale datasets for displaying overall cluster structure, these methods simply do not scale to large data with millions of edges nor do they account for attributed graphs (scalar values). Bezerianos et al. [23] and Wattenberg [24] propose interactive visual system to let users explore networks with multiple node or edge attributes. However, these methods do not consider hierarchical relationships among components-of-interest and graph attributes simultaneously.

Sariyuce et al. [25] proposed (r, s)-nucleus to denote a dense subgraph comprised by cliques, and use forest of nucleus to represent hierarchical structure of a graph. The difference is, their definition of (r, s)-nucleus focuses on density of subgraph, while our definition of ***maximal α-connected components*** focuses on relation between scalar values and graph topology, and not just limited to nucleus motifs. Moreover, their effort does not consider visualization as an objective – a primary focus of our effort. Some visualization methods (such as LaNet-vi [6]) are proposed to visualize K-Cores, but they are not general enough to handle other vertex/edge attributes. Martin et al. [26] proposed OpenOrd to visualize large scale graphs in multi-level way, but do not effectively highlight components-of-interest. Both LaNet-vi and OpenOrd share some of our objectives w.r.t network visualization, and we compare and contrast with these efforts in Section IV.

In summary, a major difference between our work and prior art is that we propose to analyze the graph through the hierarchical structure (scalar tree) induced by the ***maximal α-connected components***. The benefit is that it *naturally* encodes the relationship between clustering structure and scalar values – it highlights how scalar values evolve from high values to low values over the graph. This is particularly useful for a data scientist who wishes to understand how a community is expanded from its core members to peripheral members(see Figure 8)[27], [28], [29]. Furthermore, based on different attributes, the ***maximal α-connected component*** can represent different subgraph patterns, such as K-core, K-truss, subcommunity, which has attracted much interest within the database community [1], [5], [16], to reveal the topological relationship (containment, connection) among components of interest (e.g. K-Cores, K-Trusses, communities). We examine these issues next.

IV. User Study

In this section we briefly report on a user study which evaluates the effectiveness of our approach. Participants were recruited under OSU IRB (2015B0249) protocol, and follows recommendations of the Nielsen Norman Group (NNG). Ten participants, a sufficient number per NNG recommendations, were recruited, for *each* of the following tasks (no overlap). Each task was designed to require participants' to solve a real world problem using multiple visualization platforms. Participants were recruited from across the university campus, and were largely drawn from quantitatively inclined specializations, due to the nature of the three tasks and as a representation of the domain scientists that may use such tools (e.g. Mathematical and Physical Sciences, Psychology, Engineering, Medicine, Finance). Each participant was given a training tutorial before their prescribed task and were allowed to familiarize themselves with the functionality of the three tools used in the study.

A. Tasks

Task 1: *Identify the densest K-Core in the graph.*
The users participating in the first task were asked to identify the densest K-Core in each of three graphs (GrQc, PPI, DBLP). For each graph, we pre-computed the K-Core value (KC(v)) of each vertex, allowing participants to visualize each graph using following three different visualization methods:
(1) Our terrain visualization using K-Core value as scalar value.
(2) LaNet-vi which is a K-Core visualization tool [6].
(3) OpenOrd which is a multilevel graph layout method [26] and uses vertex's color to represent its K-Core value.
Task 1 is meaningful to data mining and network science researchers who want to explore K-Cores in a graph. The densest K-Core usually indicates a significant group of closely related nodes in the graph, such as an important community in a social network.

Task 2: *Identify the second densest K-Core in the graph that are not connected to the densest K-Core.*
In this task we provided participants with information on the densest K-Core in each dataset, and asked them to find the next densest K-core which is disconnected from the densest K-core. This is a slightly nuanced, and more complicated variant of Task 1. We note that when identifying the second K-Core for analysis simply choosing the second densest K-Core could be meaningless, because it might be heavily overlapped with the densest K-Core, and the two K-Cores are actually the same group of closely related nodes. It is more meaningful to identify the densest K-Core in the graph that are not connected to the previously detected one, because such a K-Core would indicate a separate module-of-interest.

Task 3: *The third task was the most complicated. Participants were asked to solve a problem involving multiple scalar field visualization on a single (Astro) dataset.*
For the terrain visualization method, we provided the betweenness centrality of each vertex as the first scalar value to generate the terrain, and provided the degree centrality of each vertex as the second scalar value to color the terrain (Figure 13(a)). In the visualization generated by OpenOrd, vertex color was used to indicate betweenness centrality, and vertex size to indicate degree centrality (Figure 13(b)). We ask the ten participants to determine whether the betweenness centrality and degree centrality are positively or negatively correlated in the Astro dataset. In Task 3 we do not compare with LaNet-vi, because it is specifically designed for visualizing K-Cores, and is not trivially adaptable to visualize two centralities.

B. Results

For Task 1 and Task 2, we list the average completion time and accuracy of all users in Table IV and Table V. All the pictures generated by the three visualization methods on the three datasets are from Figure 12(a) to Figure 12(i). In each picture we label the K-Core to be identified in Task 1/Task 2 as K1/K2.

In Table IV, we can see that all users successfully finished Task 1 by using terrain visualization. Two users failed by using LaNet-vi on DBLP dataset, and two users failed by using OpenOrd on the PPI dataset. The reason is that the densest

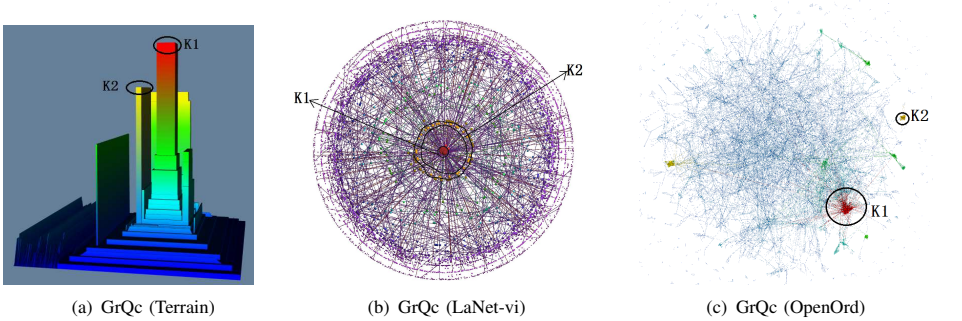

(a) GrQc (Terrain)

(b) GrQc (LaNet-vi)

(c) GrQc (OpenOrd)

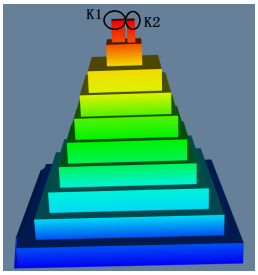

(d) PPI (Terrain)

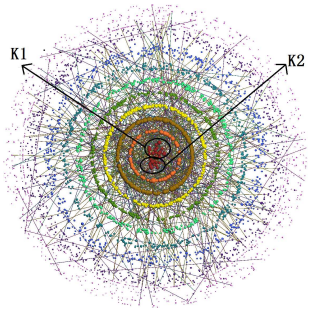

(e) PPI (LaNet-vi)

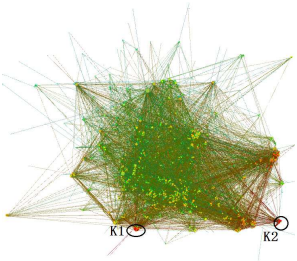

(f) PPI (OpenOrd)

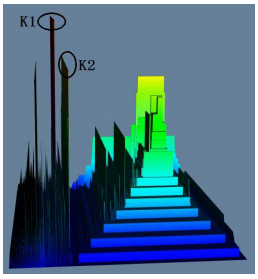

(g) DBLP (Terrain)

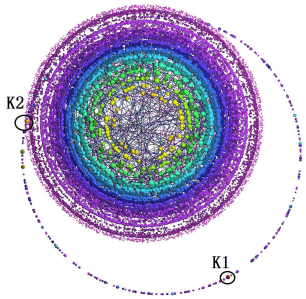

(h) DBLP (LaNet-vi)

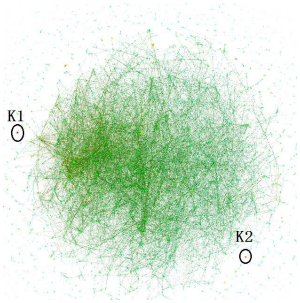

(i) DBLP (OpenOrd)

Fig. 12. Visualizations of Single Scalar Field (GrQc, PPI and DBLP)

TABLE IV. AVERAGE TIME(SEC) AND ACCURACY OF TASK 1

	Terrain		LaNet-vi		OpenOrd	
Dataset	accuracy	time	accuracy	time	accuracy	time
GrQc	1	2.6	1	6.7	1	7.6
PPI	1	4.9	1	5.3	0.8	10.7
DBLP	1	4.6	0.8	6.6	1	10.9

TABLE V. AVERAGE TIME(SEC) AND ACCURACY OF TASK 2

	Terrain		LaNet-vi		OpenOrd	
Dataset	accuracy	time	accuracy	time	accuracy	time
GrQc	1	3.6	0.7	10.3	0.8	8.2
PPI	1	4.1	0.2	7.7	0.7	11.6
DBLP	1	5.1	0.8	8.5	0.9	9.8

TABLE VI. AVERAGE TIME(SEC) AND ACCURACY OF TASK 3

	Terrain		OpenOrd	
Dataset	accuracy	time	accuracy	time
Astro	0.9	9.1	0.7	11.9

K-Core in these two visualizations are small and not obvious, and they did not notice it or they couldn't correctly identify the K-Core value through the color. Anecdotally we observed that although they zoomed-in the two visualizations to see a larger picture, they lost the full context and could only see portion of the picture, leading them to choose incorrect densest K-Cores.

Table V shows that all users successfully finished Task 2 by terrain visualization, while some users failed by using LaNet-vi and OpenOrd. One reason is Task 2 requires users to understand the connectivity to the densest K-Core (to avoid finding one with significant overlap), the LaNet-vi and OpenOrd both draw edges to indicate connections, and users need to check the edges carefully to determine whether two K-Cores are connected, it is time consuming and led to mistakes being made. In both tasks, we can see from Table IV and Table V that users spent least time on terrain visualization.

The results of Task 3 is presented in the Table VI, and the visualization pictures are in Figure 13(a) and Figure 13(b). The result shows that users achieved higher accuracy and spent less time on terrain visualization. In the visualization generated by OpenOrd, some nodes are blocked by other nodes, which caused users to make the incorrect decision. One participant also pointed out that it is easier to identify the outlier areas (circle in Figure 13(a)) in the terrain visualization than in the OpenOrd visualization. Anecdotally,several (more than 10)

users found the ability to rotate, as well as the linked 2D-Maps, a useful feature in the Terrain Visualization strategy.

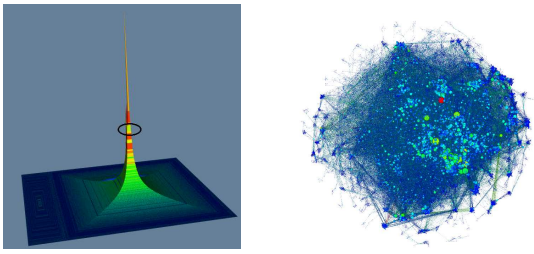

(a) Astro (Terrain) (b) Astro (OpenOrd)

Fig. 13. Visualizations of Task 3 in User Study

V. Conclusion

In this paper, we propose a visualization method to analyze scalar graphs. We demonstrate that our visualization method can reveal important information such as the overall distribution of attribute values over a graph, while simultaneously highlighting components-of-interest (dense subgraphs, social communities, etc.). It can also be extended to analyze relationship between multiple attributes. We evaluate our system on a range of real-world data and demonstrated its effectiveness and scalability on a range of data science tasks. One interesting avenue of ongoing research is to incorporate these ideas within a database – where one can view the result of a query as an attributed graph (attributed by similarities among query result tuples). We envision that such visualizations can aid and abet users in creating more natural interfaces to interrogate and understand their data.

References

[1] N. Wang, S. Parthasarathy, K.-L. Tan, and A. K. H. Tung, "Csv: Visualizing and mining cohesive subgraphs," *ACM SIGMOD*, 2008.

[2] Z. Xu, Y. Ke, Y. Wang, H. Cheng, and J. Cheng, "A model-based approach to attributed graph clustering," *ACM SIGMOD*, 2012.

[3] N. Wang, J. Zhang, K.-L. Tan, and A. K. H. Tung, "On triangulation-based dense neighborhood graph discovery," *VLDB*, 2010.

[4] Y. Zhang and S. Parthasarathy, "Extracting analyzing and visualizing triangle k-core motifs within networks," *IEEE ICDE*, 2012.

[5] V. Batagelj and M. Zaversnik, "An O(m) Algorithm for Cores Decomposition of Networks," *CoRR, arXiv.org/cs.DS/0310049*, 2003.

[6] I. Alvarez-Hamelin, L. Dall'Asta, A. Barrat, and A. Vespignani, "k-core decomposition: A tool for the analysis of large scale internet graphs." *arXiv:cs.NI/0511007*, 2005.

[7] R. Wenger, *Isosurfaces: Geometry, Topology, and Algorithms*. CRC Press, 2013.

[8] N. Sauber, H. Theisel, and H.-P. Seidel, "Multifield-graphs: An approach to visualizing correlations in multifield scalar data," *IEEE TVCG*, pp. 917–924, 2006.

[9] G. Weber, P.-T. Bremer, and V. Pascucci, "Topological landscapes: A terrain metaphor for scientific data," *IEEE TVCG*, pp. 1416–1423, 2007.

[10] W. Harvey and Y. Wang, "Generating and exploring a collection of topological landscapes for visualization of scalar-valued functions," *Eurographics*, 2010.

[11] D. Auber, Y. Chiricota, F. Jourdan, and G. Melancon, "Multiscale visualization of small world networks," *Proceedings of the Ninth Annual IEEE Conference on Information Visualization*, pp. 75–81, 2003.

[12] M. Freire, C. Plaisant, B. Shneiderman, and J. Golbeck, "Manynets: An interface for multiple network analysis and visualization," *SIGCHI*, pp. 213–222, 2010.

[13] Y. Zhou, H. Cheng, and J. X. Yu, "Graph clustering based on structural/attribute similarities," *VLDB*, 2009.

[14] J. Yang and J. Leskovec, "Overlapping community detection at scale: A nonnegative matrix factorization approach," *ACM WSDM*, 2013.

[15] H. Carr, J. Snoeyink, and U. Axen, "Computing contour trees in all dimensions," *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*.

[16] J. Wang and J. Cheng, "Truss decomposition in massive networks," *VLDB*, 2012.

[17] J. Cohen, "Trusses: Cohesive subgraphs for social network analysis," *Technical report, National Security Agency*, 2008.

[18] T. von Landesberger, A. Kuijper, T. Schreck, J. Kohlhammer, J. van Wijk, J.-D. Fekete, and D. Fellner, "Visual analysis of large graphs: State-of-the-art and future research challenges," *Computer Graphics Forum*, 2011.

[19] M. Gronemann and M. Junger, "Drawing clustered graphs as topographic maps," *Graph Drawing*, 2012.

[20] J. C. Athenstadt, R. Gorke, M. Krug, and M. Nollenburg, "Visualizing large hierarchically clustered graphs with a landscape metaphor," *Graph Drawing*, 2012.

[21] R. van Liere and W. de Leeuw, "Graphsplatting: Visualizing graphs as continuous fields," *IEEE TVCG*, 2003.

[22] A. Telea and O. Ersoy, "Image-based edge bundles: simplified visualization of large graphs," *Eurographics*, 2010.

[23] A. Bezerianos, F. Chevalier, P. Dragicevic, N. Elmqvist, and J. D. Fekete, "Graphdice: A system for exploring multivariate social networks," *Eurographics*, 2010.

[24] M. Wattenberg, "Visual exploration of multivariate graphs," *SIGCHI*, 2006.

[25] A. E. Sariyuce, C. Seshadhri, A. Pinar, and U. V. Catalyurek, "Finding the hierarchy of dense subgraphs using nucleus decompositions," *WWW*, 2015.

[26] S. Martin, W. M. Brown, R. Klavans, and K. W. Boyack;, "Openord: an open-source toolbox for large graph layout," *Proc. SPIE*, vol. 7868, 2011.

[27] C. L. Bajaj, V. Pascucci, and D. R. Schikore, "The contour spectrum," *IEEE Visualization '97*, 1997.

[28] G. H. Weber, S. E. Dillard, H. Carr, V. Pascucci, and B. Hamann, "Topology-controlled volume rendering," *IEEE TVCG*, vol. 13, no. 2, pp. 330–341, 2007.

[29] H. Carr and J. Snoeyink, "Path seeds and flexible isosurfaces using topology for exploratory visualization," *Proceedings of the symposium on Data visualisation 2003*.

[30] J. Leskovec and A. Krevl, "SNAP Datasets: Stanford large network dataset collection," http://snap.stanford.edu/data.

[31] T. M. J. Fruchterman and E. M. Reingold, "Graph drawing by force-directed placement," *Softw. Pract. Exper.*, 1991.

[32] K. Henderson, B. Gallagher, T. Eliassi-Rad, H. Tong, S. Basu, L. Akoglu, D. Koutra, C. Faloutsos, and L. Li, "Structural role extraction and mining in large graphs," *ACM SIGKDD*, pp. 1231–1239, 2012.

[33] Y. Ruan and S. Parthasarathy, "Simultaneous detection of communities and roles from large networks," *Proceedings of the 2nd ACM Conference on Online Social Networks*, pp. 203–214, 2014.